

Package ‘forecast’

April 30, 2012

Title Forecasting functions for time series and linear models

Description Methods and tools for displaying and analysing univariate time series forecasts including exponential smoothing via state space models and automatic ARIMA modelling.

Version 3.21

Date 2012-04-26

Depends R ($\geq 2.14.0$), graphics, stats, parallel, tseries, fracdiff,zoo, Rcpp ($\geq 0.9.10$), RcppArmadillo ($\geq 0.2.35$)

LinkingTo Rcpp, RcppArmadillo

LazyData yes

ByteCompile TRUE

Author Rob J Hyndman <Rob.Hyndman@monash.edu> with contributions from Slava Razbash and Drew Schmidt

Maintainer Rob J Hyndman <Rob.Hyndman@monash.edu>

License GPL (≥ 2)

URL <http://robjhyndman.com/software/forecast/>

Repository CRAN

Date/Publication 2012-04-30 05:37:32

R topics documented:

accuracy	3
Acf	4
arfima	5
Arima	6
arima.errors	8
auto.arima	9

bats	11
BoxCox	12
BoxCox.lambda	13
croston	14
CV	16
dm.test	17
dshw	18
ets	20
fitted.Arima	22
forecast	23
forecast.Arima	24
forecast.bats	26
forecast.ets	27
forecast.HoltWinters	29
forecast.lm	30
forecast.stl	32
forecast.StructTS	34
gas	35
gold	36
logLik.ets	36
ma	37
meanf	38
monthdays	39
msts	40
na.interp	41
naive	41
ndiffs	43
plot.bats	44
plot.ets	45
plot.forecast	46
rwf	47
seasadj	49
seasonaldummy	49
seasonplot	51
ses	52
simulate.ets	53
sindexf	55
splinef	55
subset.ts	57
taylor	58
tbats	58
thetaf	60
tsdisplay	61
tslm	62
wineind	63
woolyrnq	64

accuracy

Accuracy measures for forecast model

Description

Returns range of summary measures of the forecast accuracy. If x is provided, the function measures out-of-sample forecast accuracy based on x - f . If x is not provided, the function produces in-sample accuracy measures of the forecasts based on $f["x"]$ -fitted(f). All measures are defined and discussed in Hyndman and Koehler (2006).

Usage

```
accuracy(f, x, test="all")
```

Arguments

<code>f</code>	An object of class "forecast", or a numerical vector containing forecasts. It will also work with Arima, ets and lm objects if x is omitted – in which case in-sample accuracy measures are returned.
<code>x</code>	An optional numerical vector containing actual values of the same length as object.
<code>test</code>	Indicator of which elements of x and f to test. If <code>test=="all"</code> , all elements are used. Otherwise <code>test</code> is a numeric vector containing the indices of the elements to use in the test.

Details

MASE calculation is scaled using MAE of in-sample naive forecasts for non-seasonal time series, in-sample seasonal naive forecasts for seasonal time series and in-sample mean forecasts for non-time series data.

Value

Vector giving forecast accuracy measures.

Author(s)

Rob J Hyndman

References

Hyndman, R.J. and Koehler, A.B. (2006) "Another look at measures of forecast accuracy". *International Journal of Forecasting*, **22**(4).

Examples

```
fit1 <- rwf(EuStockMarkets[1:200,1],h=100)
fit2 <- meanf(EuStockMarkets[1:200,1],h=100)
accuracy(fit1)
accuracy(fit2)
accuracy(fit1,EuStockMarkets[201:300,1])
accuracy(fit2,EuStockMarkets[201:300,1])
plot(fit1)
lines(EuStockMarkets[1:300,1])
```

Acf

(Partial) Autocorrelation Function Estimation

Description

Largely wrappers for the [acf](#) function in the stats package. The main difference is that Acf does not plot a spike at lag 0 (which is redundant). Pacf is included for consistency.

Usage

```
Acf(x, lag.max=NULL, type=c("correlation", "partial"),
    plot=TRUE, main=NULL, ...)
Pacf(x, main=NULL, ...)
```

Arguments

x	a univariate time series
lag.max	maximum lag at which to calculate the acf. Default is $10 \cdot \log_{10}(N/m)$ where N is the number of observations and m the number of series. Will be automatically limited to one less than the number of observations in the series.
type	character string giving the type of acf to be computed. Allowed values are "correlation" (the default) or "partial".
plot	logical. If TRUE (the default) the acf is plotted.
main	Title for plot
...	Additional arguments passed to acf .

Details

See the [acf](#) function in the stats package.

Value

See the [acf](#) function in the stats package.

Author(s)

Rob J Hyndman

See Also[acf](#)**Examples**

```
Acf(wineind)
Pacf(wineind)
```

`arfima`*Fit a fractionally differenced ARFIMA model*

Description

An ARFIMA(p,d,q) model is selected and estimated automatically using the Hyndman-Khandakar (2008) algorithm to select p and q and the Haslett and Raftery (1989) algorithm to estimate the parameters including d.

Usage

```
arfima(x, drange=c(0, 0.5), estim=c("mle","ls"), lambda=NULL, ...)
```

Arguments

<code>x</code>	a univariate time series (numeric vector).
<code>drange</code>	Allowable values of d to be considered. Default of <code>c(0, 0.5)</code> ensures a stationary model is returned.
<code>estim</code>	If <code>estim=="ls"</code> , then the ARMA parameters are calculated using the Haslett-Raftery algorithm. If <code>estim=="mle"</code> , then the ARMA parameters are calculated using full MLE via the arfima function.
<code>lambda</code>	Box-Cox transformation parameter. Ignored if NULL. Otherwise, data transformed before model is estimated.
<code>...</code>	Other arguments passed to auto.arima when selecting p and q.

Details

This function combines [fracdiff](#) and [auto.arima](#) to automatically select and estimate an ARFIMA model. The fractional differencing parameter is chosen first assuming an ARFIMA(2,d,0) model. Then the data are fractionally differenced using the estimated d and an ARMA model is selected for the resulting time series using [auto.arima](#). Finally, the full ARFIMA(p,d,q) model is re-estimated using [fracdiff](#). If `estim=="mle"`, the ARMA coefficients are refined using [arfima](#).

Value

A list object of S3 class "fracdiff", which is described in the [fracdiff](#) documentation. A few additional objects are added to the list including `x` (the original time series), and the residuals and fitted values.

Author(s)

Rob J Hyndman and Farah Yasmeeen

References

J. Haslett and A. E. Raftery (1989) Space-time Modelling with Long-memory Dependence: Assessing Ireland's Wind Power Resource (with discussion); *Applied Statistics* **38**, 1-50.

Hyndman, R.J. and Khandakar, Y. (2008) "Automatic time series forecasting: The forecast package for R", *Journal of Statistical Software*, **26**(3).

See Also

[fracdiff](#), [auto.arima](#), [forecast.fracdiff](#).

Examples

```
x <- fracdiff.sim( 100, ma=-.4, d=.3)$series
fit <- arfima(x)
tsdisplay(residuals(fit))
```

Arima

Fit ARIMA model to univariate time series

Description

Largely a wrapper for the [arima](#) function in the stats package. The main difference is that this function allows a drift term. It is also possible to take an ARIMA model from a previous call to Arima and re-apply it to the data x.

Usage

```
Arima(x, order=c(0,0,0), seasonal=list(order=c(0,0,0), period=NA),
      xreg=NULL, include.mean=TRUE, include.drift=FALSE,
      include.constant, lambda=model$lambda, transform.pars=TRUE,
      fixed=NULL, init=NULL, method=c("CSS-ML", "ML", "CSS"), n.cond,
      optim.control=list(), kappa=1e6, model=NULL)
```

Arguments

x	a univariate time series
order	A specification of the non-seasonal part of the ARIMA model: the three components (p, d, q) are the AR order, the degree of differencing, and the MA order.
seasonal	A specification of the seasonal part of the ARIMA model, plus the period (which defaults to frequency(x)). This should be a list with components order and period, but a specification of just a numeric vector of length 3 will be turned into a suitable list with the specification as the order.

<code>xreg</code>	Optionally, a vector or matrix of external regressors, which must have the same number of rows as <code>x</code> .
<code>include.mean</code>	Should the ARIMA model include a mean term? The default is TRUE for undifferenced series, FALSE for differenced ones (where a mean would not affect the fit nor predictions).
<code>include.drift</code>	Should the ARIMA model include a linear drift term? (i.e., a linear regression with ARIMA errors is fitted.) The default is FALSE.
<code>include.constant</code>	If TRUE, then <code>include.mean</code> is set to be TRUE for undifferenced series and <code>include.drift</code> is set to be TRUE for differenced series. Note that if there is more than one difference taken, no constant is included regardless of the value of this argument. This is deliberate as otherwise quadratic and higher order polynomial trends would be induced.
<code>lambda</code>	Box-Cox transformation parameter. Ignored if NULL. Otherwise, data transformed before model is estimated.
<code>transform.pars</code>	Logical. If true, the AR parameters are transformed to ensure that they remain in the region of stationarity. Not used for <code>method="CSS"</code> .
<code>fixed</code>	optional numeric vector of the same length as the total number of parameters. If supplied, only NA entries in <code>fixed</code> will be varied. <code>transform.pars=TRUE</code> will be overridden (with a warning) if any AR parameters are fixed. It may be wise to set <code>transform.pars=FALSE</code> when fixing MA parameters, especially near non-invertibility.
<code>init</code>	optional numeric vector of initial parameter values. Missing values will be filled in, by zeroes except for regression coefficients. Values already specified in <code>fixed</code> will be ignored.
<code>method</code>	Fitting method: maximum likelihood or minimize conditional sum-of-squares. The default (unless there are missing values) is to use conditional-sum-of-squares to find starting values, then maximum likelihood.
<code>n.cond</code>	Only used if fitting by conditional-sum-of-squares: the number of initial observations to ignore. It will be ignored if less than the maximum lag of an AR term.
<code>optim.control</code>	List of control parameters for <code>optim</code> .
<code>kappa</code>	the prior variance (as a multiple of the innovations variance) for the past observations in a differenced model. Do not reduce this.
<code>model</code>	Output from a previous call to <code>Arima</code> . If <code>model</code> is passed, this same model is fitted to <code>x</code> without re-estimating any parameters.

Details

See the [arima](#) function in the stats package.

Value

See the [arima](#) function in the stats package. The additional objects returned are

<code>x</code>	The time series data
<code>xreg</code>	The regressors used in fitting (when relevant).

Author(s)

Rob J Hyndman

See Also[arima](#)**Examples**

```
fit <- Arima(WWWusage,order=c(3,1,0))
plot(forecast(fit,h=20))

# Fit model to first few years of AirPassengers data
air.model <- Arima(window(AirPassengers,end=1956+11/12),order=c(0,1,1),
                  seasonal=list(order=c(0,1,1),period=12),lambda=0)
plot(forecast(air.model,h=48))
lines(AirPassengers)

# Apply fitted model to later data
air.model2 <- Arima(window(AirPassengers,start=1957),model=air.model)

# Forecast accuracy measures on the log scale.
# in-sample one-step forecasts.
accuracy(air.model)
# out-of-sample one-step forecasts.
accuracy(air.model2)
# out-of-sample multi-step forecasts
accuracy(forecast(air.model,h=48,lambda=NULL),
        log(window(AirPassengers,start=1957)))
```

arima.errors

ARIMA errors

Description

Returns original time series after adjusting for regression variables. These are not the same as the residuals. If there are no regression variables in the ARIMA model, then the errors will be identical to the original series. If there are regression variables in the ARIMA model, then the errors will be equal to the original series minus the effect of the regression variables, but leaving in the serial correlation that is modelled with the AR and MA terms. If you want the "residuals", then use `residuals(z)`.

Usage

```
arima.errors(z)
```

Arguments

`z` Fitted ARIMA model from [arima](#)

Value

A time series containing the "errors".

Author(s)

Rob J Hyndman

See Also

[arima](#), [residuals](#)

Examples

```
ukdeaths.fit <- Arima(UKDriverDeaths,c(1,0,1),c(0,1,1),xreg=Seatbelts[, "law"])
ukdeaths.errors <- arima.errors(ukdeaths.fit)
par(mfrow=c(2,1))
plot(UKDriverDeaths)
plot(ukdeaths.errors)
```

auto.arima

Fit best ARIMA model to univariate time series

Description

Returns best ARIMA model according to either AIC, AICc or BIC value. The function conducts a search over possible model within the order constraints provided.

Usage

```
auto.arima(x, d=NA, D=NA, max.p=5, max.q=5,
  max.P=2, max.Q=2, max.order=5, start.p=2, start.q=2,
  start.P=1, start.Q=1, stationary=FALSE,
  ic=c("aicc", "aic", "bic"), stepwise=TRUE, trace=FALSE,
  approximation=(length(x)>100 | frequency(x)>12), xreg=NULL,
  test=c("kpss", "adf", "pp"), seasonal.test=c("ocsb", "ch"),
  allowdrift=TRUE, lambda=NULL, parallel=FALSE, num.cores=NULL)
```

Arguments

x	a univariate time series
d	Order of first-differencing. If missing, will choose a value based on KPSS test.
D	Order of seasonal-differencing. If missing, will choose a value based on CH test.
max.p	Maximum value of p
max.q	Maximum value of q
max.P	Maximum value of P

max.Q	Maximum value of Q
max.order	Maximum value of $p+q+P+Q$ if model selection is not stepwise.
start.p	Starting value of p in stepwise procedure.
start.q	Starting value of q in stepwise procedure.
start.P	Starting value of P in stepwise procedure.
start.Q	Starting value of Q in stepwise procedure.
stationary	If TRUE, restricts search to stationary models.
ic	Information criterion to be used in model selection.
stepwise	If TRUE, will do stepwise selection (faster). Otherwise, it searches over all models. Non-stepwise selection can be very slow, especially for seasonal models.
trace	If TRUE, the list of ARIMA models considered will be reported.
approximation	If TRUE, estimation is via conditional sums of squares and the information criteria used for model selection are approximated. The final model is still computed using maximum likelihood estimation. Approximation should be used for long time series or a high seasonal period to avoid excessive computation times.
xreg	Optionally, a vector or matrix of external regressors, which must have the same number of rows as x.
test	Type of unit root test to use. See ndiffs for details.
seasonal.test	This determines which seasonal unit root test is used. See nsdiffs for details.
allowdrift	If TRUE, models with drift terms are considered.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, data transformed before model is estimated.
parallel	If TRUE and <code>stepwise = FALSE</code> , then the specification search is done in parallel. This can give a significant speedup on multicore machines.
num.cores	Allows the user to specify the amount of parallel processes to be used if <code>parallel = TRUE</code> and <code>stepwise = FALSE</code> . If NULL, then the number of logical cores is automatically detected.

Details

Non-stepwise selection can be slow, especially for seasonal data. Stepwise algorithm outlined in Hyndman and Khandakar (2008) except that the default method for selecting seasonal differences is now the OCSB test rather than the Canova-Hansen test.

Value

Same as for [arima](#)

Author(s)

Rob J Hyndman

References

Hyndman, R.J. and Khandakar, Y. (2008) "Automatic time series forecasting: The forecast package for R", *Journal of Statistical Software*, **26**(3).

See Also

[Arima](#)

Examples

```
fit <- auto.arima(WWWusage)
plot(forecast(fit,h=20))
```

bats	<i>BATS model (Exponential smoothing state space model with Box-Cox transformation, ARMA errors, Trend and Seasonal components)</i>
------	---

Description

Fits a BATS model applied to y , as described in De Livera, Hyndman & Snyder (2011). Parallel processing is used by default to speed up the computations.

Usage

```
bats(y, use.box.cox=NULL, use.trend=NULL, use.damped.trend=NULL,
     seasonal.periods=NULL, use.arma.errors=TRUE, use.parallel=TRUE, num.cores=NULL, ...)
```

Arguments

<code>y</code>	The time series to be forecast. Can be numeric, msts or ts. Only univariate time series are supported.
<code>use.box.cox</code>	TRUE/FALSE indicates whether to use the Box-Cox transformation or not. If NULL then both are tried and the best fit is selected by AIC.
<code>use.trend</code>	TRUE/FALSE indicates whether to include a trend or not. If NULL then both are tried and the best fit is selected by AIC.
<code>use.damped.trend</code>	TRUE/FALSE indicates whether to include a damping parameter in the trend or not. If NULL then both are tried and the best fit is selected by AIC.
<code>seasonal.periods</code>	If y is a numeric then seasonal periods can be specified with this parameter.
<code>use.arma.errors</code>	TRUE/FALSE indicates whether to include ARMA errors or not. If TRUE the best fit is selected by AIC. If FALSE then the selection algorithm does not consider ARMA errors.
<code>use.parallel</code>	TRUE/FALSE indicates whether or not to use parallel processing.

num.cores The number of parallel processes to be used if using parallel processing. If NULL then the number of logical cores is detected.

... Additional parameters to be passed to `auto.arima` when choose an ARMA(p, q) model for the errors.

Value

An object of class "bats". The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `bats` and associated functions.

Author(s)

Slava Razbash and Rob J Hyndman

References

De Livera, A.M., Hyndman, R.J., & Snyder, R. D. (2011), Forecasting time series with complex seasonal patterns using exponential smoothing, *Journal of the American Statistical Association*, **106**(496), 1513-1527.

Examples

```
fit <- bats(USAccDeaths)
plot(forecast(fit))
## Not run:
taylor.fit <- bats(taylor)
plot(forecast(taylor.fit))

## End(Not run)
```

BoxCox

Box Cox Transformation

Description

`BoxCox()` returns a transformation of the input variable using a Box-Cox transformation. `InvBoxCox()` reverses the transformation.

Usage

```
BoxCox(x, lambda)
InvBoxCox(x, lambda)
```

Arguments

x a numeric vector or time series

lambda transformation parameter

Details

The Box-Cox transformation is given by

$$f_{\lambda}(x) = \frac{x^{\lambda} - 1}{\lambda}$$

if $\lambda \neq 0$. For $\lambda = 0$,

$$f_0(x) = \log(x)$$

Value

a numeric vector of the same length as x .

Author(s)

Rob J Hyndman

References

Box, G. E. P. and Cox, D. R. (1964) An analysis of transformations. *JRSS B* **26** 211–246.

See Also

[BoxCox.lambda](#)

Examples

```
lambda <- BoxCox.lambda(lynx)
lynx.fit <- ar(BoxCox(lynx, lambda))
plot(forecast(lynx.fit, h=20, lambda=lambda))
```

BoxCox.lambda

Automatic selection of Box Cox transformation parameter

Description

If `method=="guerrero"`, Guerrero's (1993) method is used, where `lambda` minimizes the coefficient of variation for subseries of x .

If `method=="loglik"`, the value of `lambda` is chosen to maximize the profile log likelihood of a linear model fitted to x . For non-seasonal data, a linear time trend is fitted while for seasonal data, a linear time trend with seasonal dummy variables is used.

Usage

```
BoxCox.lambda(x, method=c("guerrero", "loglik"), lower=-1, upper=2)
```

Arguments

x	a numeric vector or time series
method	Choose method to be used in calculating lambda.
lower	Lower limit for possible lambda values.
upper	Upper limit for possible lambda values.

Value

a number indicating the Box-Cox transformation parameter.

Author(s)

Leanne Chhay and Rob J Hyndman

References

- Box, G. E. P. and Cox, D. R. (1964) An analysis of transformations. *JRSS B* **26** 211–246.
- Guerrero, V.M. (1993) Time-series analysis supported by power transformations. *Journal of Forecasting*, **12**, 37–48.

See Also

[BoxCox](#)

Examples

```
lambda <- BoxCox.lambda(AirPassengers, lower=0)
air.fit <- Arima(AirPassengers, order=c(0,1,1),
                seasonal=list(order=c(0,1,1), period=12), lambda=lambda)
plot(forecast(air.fit))
```

croston

Forecasts for intermittent demand using Croston's method

Description

Returns forecasts and other information for Croston's forecasts applied to x.

Usage

```
croston(x, h=10, alpha=0.1)
```

Arguments

x	a numeric vector or time series
h	Number of periods for forecasting.
alpha	Value of alpha. Default value is 0.1.

Details

Based on Croston's (1972) method for intermittent demand forecasting, also described in Shenstone and Hyndman (2005). Croston's method involves using simple exponential smoothing (SES) on the non-zero elements of the time series and a separate application of SES to the times between non-zero elements of the time series. The smoothing parameters of the two applications of SES are assumed to be equal and are denoted by alpha.

Note that prediction intervals are not computed as Croston's method has no underlying stochastic model.

Value

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model. The first element gives the SES model used for non-zero demands. The second element gives the SES model used for times between non-zero demands. Both models are of class forecast.
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `croston` and associated functions.

Author(s)

Rob J Hyndman

References

Croston, J. (1972) "Forecasting and stock control for intermittent demands", *Operational Research Quarterly*, **23**(3), 289-303.

Shenstone, L., and Hyndman, R.J. (2005) "Stochastic models underlying Croston's method for intermittent demand forecasting". *Journal of Forecasting*, **24**, 389-402.

See Also

[ses](#).

Examples

```
x <- rpois(20,lambda=.3)
fcast <- croston(x)
plot(fcast)
```

CV

Cross-validation statistic

Description

Computes cross-validation statistic, AIC, corrected AIC, BIC and adjusted R² values for a linear model.

Usage

```
CV(obj)
```

Arguments

obj output from [lm](#) or [tslm](#)

Value

Numerical vector containing CV, AIC, AICc, BIC and AdjR² values.

Author(s)

Rob J Hyndman

See Also

[AIC](#)

Examples

```
y <- ts(rnorm(120,0,3) + 20*sin(2*pi*(1:120)/12), frequency=12)
fit1 <- tslm(y ~ trend + season)
fit2 <- tslm(y ~ season)
CV(fit1)
CV(fit2)
```

dm.test	<i>Diebold-Mariano test for predictive accuracy</i>
---------	---

Description

The Diebold-Mariano test compares the forecast accuracy of two forecast methods. The null hypothesis is that they have the same forecast accuracy.

Usage

```
dm.test(e1, e2, alternative=c("two.sided", "less", "greater"),
        h=1, power=2)
```

Arguments

e1	Forecast errors from method 1.
e2	Forecast errors from method 2.
alternative	a character string specifying the alternative hypothesis, must be one of "two.sided" (default), "greater" or "less". You can specify just the initial letter.
h	The forecast horizon used in calculating e1 and e2.
power	The power used in the loss function. Usually 1 or 2.

Value

A list with class "htest" containing the following components:

statistic	the value of the DM-statistic.
parameter	the forecast horizon and loss function power used in the test.
alternative	a character string describing the alternative hypothesis.
p.value	the p-value for the test.
method	a character string with the value "Diebold-Mariano Test".
data.name	a character vector giving the names of the two error series.

Author(s)

George Athanasopoulos and Rob Hyndman

References

Diebold, F.X. and Mariano, R.S. (1995) Comparing predictive accuracy. *Journal of Business and Economic Statistics*, **13**, 253-263.

Examples

```
# Test on in-sample one-step forecasts
f1 <- ets(WWWusage)
f2 <- auto.arima(WWWusage)
accuracy(f1)
accuracy(f2)
dm.test(residuals(f1),residuals(f2),h=1)

# Test on out-of-sample one-step forecasts
f1 <- ets(WWWusage[1:80])
f2 <- auto.arima(WWWusage[1:80])
f1.out <- ets(WWWusage[81:100],model=f1)
f2.out <- Arima(WWWusage[81:100],model=f2)
accuracy(f1.out)
accuracy(f2.out)
dm.test(residuals(f1.out),residuals(f2.out),h=1)
```

dshw

*Double-Seasonal Holt-Winters Forecasting***Description**

Returns forecasts and prediction intervals using Taylor's (2003) Double-Seasonal Holt-Winters method.

Usage

```
dshw(y, period1, period2, h=2*max(period1,period2),
      alpha=NULL, beta=NULL, gamma=NULL, omega=NULL, phi=NULL,
      lambda=NULL, armethod=TRUE)
```

Arguments

y	a numeric vector or time series
period1	Period of the shorter seasonal period.
period2	Period of the longer seasonal period.
h	Number of periods for forecasting
alpha	Smoothing parameter for the level.
beta	Smoothing parameter for the slope.
gamma	Smoothing parameter for the first seasonal period.
omega	Smoothing parameter for the second seasonal period.
phi	Autoregressive parameter.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, data transformed before model is estimated.
armethod	If TRUE, the forecasts are adjusted using an AR(1) model for the errors.

Details

Taylor's (2003) double-seasonal Holt-Winters method uses additive trend and multiplicative seasonality, where there are two seasonal components which are multiplied together. For example, with a series of half-hourly data, one would set `period1=48` for the daily period and `period2=336` for the weekly period. The smoothing parameter notation used here is different from that in Taylor (2003); instead it matches that used in Hyndman et al (2008) and that used for the `ets` function.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `meanf`.

An object of class "forecast" is a list containing at least the following elements:

<code>model</code>	A list containing information about the fitted model
<code>method</code>	The name of the forecasting method as a character string
<code>mean</code>	Point forecasts as a time series
<code>lower</code>	Lower limits for prediction intervals
<code>upper</code>	Upper limits for prediction intervals
<code>level</code>	The confidence values associated with the prediction intervals
<code>x</code>	The original time series (either object itself or the time series used to create the model stored as object).
<code>residuals</code>	Residuals from the fitted model. That is <code>x</code> minus fitted values.
<code>fitted</code>	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

References

Taylor, J.W. (2003) Short-term electricity demand forecasting using double seasonal exponential smoothing. *Journal of the Operational Research Society*, **54**, 799-805.

Hyndman, R.J., Koehler, A.B., Ord, J.K., and Snyder, R.D. (2008) *Forecasting with exponential smoothing: the state space approach*, Springer-Verlag. <http://www.exponentialsMOOTHING.net>.

See Also

[HoltWinters](#), [ets](#).

Examples

```
## Not run:
fcast <- dshw(taylor)
plot(fcast)

## End(Not run)

t <- seq(0,5,by=1/20)
x <- exp(sin(2*pi*t)) + cos(2*pi*t*4) + rnorm(length(t),0,.1)
fit <- dshw(x,20,5)
plot(fit)
```

ets

*Exponential smoothing state space model***Description**

Returns ets model applied to y.

Usage

```
ets(y, model="ZZZ", damped=NULL, alpha=NULL, beta=NULL, gamma=NULL,
    phi=NULL, additive.only=FALSE, lambda=NULL,
    lower=c(rep(0.0001,3), 0.8), upper=c(rep(0.9999,3),0.98),
    opt.crit=c("lik","amse","mse","sigma","mae"), nmse=3,
    bounds=c("both","usual","admissible"), ic=c("aic","aicc","bic"),
    restrict=TRUE)
```

Arguments

y	a numeric vector or time series
model	Usually a three-character string identifying method using the framework terminology of Hyndman et al. (2002) and Hyndman et al. (2008). The first letter denotes the error type ("A", "M" or "Z"); the second letter denotes the trend type ("N", "A", "M" or "Z"); and the third letter denotes the season type ("N", "A", "M" or "Z"). In all cases, "N"=none, "A"=additive, "M"=multiplicative and "Z"=automatically selected. So, for example, "ANN" is simple exponential smoothing with additive errors, "MAM" is multiplicative Holt-Winters' method with multiplicative errors, and so on. It is also possible for the model to be equal to the output from a previous call to ets. In this case, the same model is fitted to y without re-estimating any parameters.
damped	If TRUE, use a damped trend (either additive or multiplicative). If NULL, both damped and non-damped trends will be tried and the best model (according to the information criterion ic) returned.
alpha	Value of alpha. If NULL, it is estimated.
beta	Value of beta. If NULL, it is estimated.

gamma	Value of gamma. If NULL, it is estimated.
phi	Value of phi. If NULL, it is estimated.
additive.only	If TRUE, will only consider additive models. Default is FALSE.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, data transformed before model is estimated. When lambda=TRUE, additive.only is set to FALSE.
lower	Lower bounds for the parameters (alpha, beta, gamma, phi)
upper	Upper bounds for the parameters (alpha, beta, gamma, phi)
opt.crit	Optimization criterion. One of "mse" (Mean Square Error), "amse" (Average MSE over first nmse forecast horizons), "sigma" (Standard deviation of residuals), "mae" (Mean of absolute residuals), or "lik" (Log-likelihood, the default).
nmse	Number of steps for average multistep MSE ($1 \leq \text{nmse} \leq 10$).
bounds	Type of parameter space to impose: "usual" indicates all parameters must lie between specified lower and upper bounds; "admissible" indicates parameters must lie in the admissible space; "both" (default) takes the intersection of these regions.
ic	Information criterion to be used in model selection.
restrict	If TRUE, the models with infinite variance will not be allowed.

Details

Based on the classification of methods as described in Hyndman et al (2008).

The methodology is fully automatic. The only required argument for ets is the time series. The model is chosen automatically if not specified. This methodology performed extremely well on the M3-competition data. (See Hyndman, et al, 2002, below.)

Value

An object of class "ets".

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by ets and associated functions.

Author(s)

Rob J Hyndman

References

Hyndman, R.J., Koehler, A.B., Snyder, R.D., and Grose, S. (2002) "A state space framework for automatic forecasting using exponential smoothing methods", *International J. Forecasting*, **18**(3), 439–454.

Hyndman, R.J., Akram, Md., and Archibald, B. (2008) "The admissible parameter space for exponential smoothing models". *Annals of Statistical Mathematics*, **60**(2), 407–426.

Hyndman, R.J., Koehler, A.B., Ord, J.K., and Snyder, R.D. (2008) *Forecasting with exponential smoothing: the state space approach*, Springer-Verlag. <http://www.exponentialsMOOTHING.net>.

See Also

[HoltWinters](#), [rwf](#), [arima](#).

Examples

```
fit <- ets(USAccDeaths)
plot(forecast(fit))
```

fitted.Arima

One-step in-sample forecasts using ARIMA models

Description

Returns one-step forecasts for the data used in fitting the ARIMA model.

Usage

```
## S3 method for class 'Arima'
fitted(object,...)
```

Arguments

`object` An object of class "Arima". Usually the result of a call to [arima](#).
`...` Other arguments.

Value

An time series of the one-step forecasts.

Author(s)

Rob J Hyndman

See Also

[forecast.Arima](#).

Examples

```
fit <- Arima(WWWusage,c(3,1,0))
plot(WWWusage)
lines(fitted(fit),col=2)
```

forecast	<i>Forecasting time series</i>
----------	--------------------------------

Description

`forecast` is a generic function for forecasting from time series or time series models. The function invokes particular *methods* which depend on the class of the first argument.

For example, the function `forecast.Arima` makes forecasts based on the results produced by `arima`.

The function `forecast.ts` makes forecasts using `ets` models (if the data are non-seasonal or the seasonal period is 12 or less) or `stlf` (if the seasonal period is 13 or more).

Usage

```
forecast(object, ...)
## S3 method for class 'ts'
forecast(object, h, level=c(80,95), fan=FALSE, ...)
```

Arguments

<code>object</code>	a time series or time series model for which forecasts are required
<code>h</code>	Number of periods for forecasting
<code>level</code>	Confidence level for prediction intervals.
<code>fan</code>	If TRUE, level is set to <code>seq(50,99,by=1)</code> . This is suitable for fan plots.
<code>...</code>	Additional arguments affecting the forecasts produced. <code>forecast.ts</code> passes these to <code>forecast.ets</code> or <code>stlf</code> depending on the frequency of the time series.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract various useful features of the value returned by `forecast$model`.

An object of class "forecast" is a list containing at least the following elements:

<code>model</code>	A list containing information about the fitted model
<code>method</code>	The name of the forecasting method as a character string
<code>mean</code>	Point forecasts as a time series
<code>lower</code>	Lower limits for prediction intervals
<code>upper</code>	Upper limits for prediction intervals
<code>level</code>	The confidence values associated with the prediction intervals

x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

See Also

Other functions which return objects of class "forecast" are [forecast.ets](#), [forecast.Arima](#), [forecast.HoltWinters](#), [forecast.StructTS](#), [meanf](#), [rwf](#), [splinef](#), [thetaf](#), [croston](#), [ses](#), [holt](#), [hw](#).

forecast.Arima	<i>Forecasting using ARIMA or ARFIMA models</i>
----------------	---

Description

Returns forecasts and other information for univariate ARIMA models.

Usage

```
## S3 method for class 'Arima'
forecast(object, h=ifelse(object$arma[5]>1,2*object$arma[5],10),
         level=c(80,95), fan=FALSE, xreg=NULL, lambda=object$lambda, ...)
## S3 method for class 'ar'
forecast(object, h=10, level=c(80,95), fan=FALSE, lambda=NULL, ...)
## S3 method for class 'fracdiff'
forecast(object, h=10, level=c(80,95), fan=FALSE, lambda=object$lambda, ...)
```

Arguments

object	An object of class "Arima", "ar" or "fracdiff". Usually the result of a call to arima , auto.arima , ar , arfima or fracdiff .
h	Number of periods for forecasting. If xreg is used, h is ignored and the number of forecast periods is set to the number of rows of xreg.
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to seq(50,99,by=1). This is suitable for fan plots.
xreg	Future values of an regression variables (for class Arima objects only).
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.
...	Other arguments.

Details

For Arima or ar objects, the function calls `predict.Arima` or `predict.ar` and constructs an object of class "forecast" from the results. For fracdiff objects, the calculations are all done within `forecast.fracdiff` using the equations given by Peiris and Perera (1988).

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.Arima`.

An object of class "forecast" is a list containing at least the following elements:

<code>model</code>	A list containing information about the fitted model
<code>method</code>	The name of the forecasting method as a character string
<code>mean</code>	Point forecasts as a time series
<code>lower</code>	Lower limits for prediction intervals
<code>upper</code>	Upper limits for prediction intervals
<code>level</code>	The confidence values associated with the prediction intervals
<code>x</code>	The original time series (either object itself or the time series used to create the model stored as object).
<code>residuals</code>	Residuals from the fitted model. That is x minus fitted values.
<code>fitted</code>	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

References

Peiris, M. & Perera, B. (1988), On prediction with fractionally differenced ARIMA models, *Journal of Time Series Analysis*, **9**(3), 215-220.

See Also

`predict.Arima`, `predict.ar`, `auto.arima`, `Arima`, `arima`, `ar`, `arfima`.

Examples

```
fit <- Arima(WWWusage,c(3,1,0))
plot(forecast(fit))

x <- fracdiff.sim( 100, ma=-.4, d=.3)$series
fit <- arfima(x)
plot(forecast(fit,h=30))
```

forecast.bats

*Forecasting using BATS and TBATS models***Description**

Forecasts h steps ahead with a BATS model. Prediction intervals are also produced.

Usage

```
## S3 method for class 'bats'
forecast(object, h, level=c(80,95), fan=FALSE, ...)
## S3 method for class 'tbats'
forecast(object, h, level=c(80,95), fan=FALSE, ...)
```

Arguments

object	An object of class "bats". Usually the result of a call to bats .
h	Number of periods for forecasting. Default value is twice the largest seasonal period (for seasonal data) or ten (for non-seasonal data).
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to seq(50, 99, by=1). This is suitable for fan plots.
...	Other arguments, currently ignored.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.bats`.

An object of class "forecast" is a list containing at least the following elements:

model	A copy of the bats object
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Slava Razbash and Rob J Hyndman

References

De Livera, A.M., Hyndman, R.J., & Snyder, R. D. (2011), Forecasting time series with complex seasonal patterns using exponential smoothing, *Journal of the American Statistical Association*, **106**(496), 1513-1527.

See Also

[bats](#), [tbats](#), [forecast.ets](#).

Examples

```
fit <- bats(USAccDeaths)
plot(forecast(fit))
## Not run:
taylor.fit <- bats(taylor)
plot(forecast(taylor.fit))

## End(Not run)
```

forecast.ets

Forecasting using ETS models

Description

Returns forecasts and other information for univariate ETS models.

Usage

```
## S3 method for class 'ets'
forecast(object, h=ifelse(object$m>1, 2*object$m, 10),
         level=c(80,95), fan=FALSE, simulate=FALSE, bootstrap=FALSE,
         npaths=5000, PI=TRUE, lambda=object$lambda, ...)
```

Arguments

object	An object of class "ets". Usually the result of a call to ets .
h	Number of periods for forecasting
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to seq(50,99,by=1). This is suitable for fan plots.
simulate	If TRUE, prediction intervals produced by simulation rather than using analytic formulae.
bootstrap	If TRUE, and if simulate=TRUE, then simulation uses resampled errors rather than normally distributed errors.

npaths	Number of sample paths used in computing simulated prediction intervals.
PI	If TRUE, prediction intervals are produced, otherwise only point forecasts are calculated. If PI is FALSE, then level, fan, simulate, bootstrap and npaths are all ignored.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.
...	Other arguments.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.ets`.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

See Also

[ets](#), [ses](#), [holt](#), [hw](#).

Examples

```
fit <- ets(USAccDeaths)
plot(forecast(fit,h=48))
```

forecast.HoltWinters *Forecasting using Holt-Winters objects*

Description

Returns forecasts and other information for univariate Holt-Winters time series models.

Usage

```
## S3 method for class 'HoltWinters'
forecast(object, h=ifelse(frequency(object$x)>1,2*frequency(object$x),10),
         level=c(80,95),fan=FALSE,lambda=NULL,...)
```

Arguments

object	An object of class "HoltWinters". Usually the result of a call to HoltWinters .
h	Number of periods for forecasting
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to seq(50,99,by=1). This is suitable for fan plots.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.
...	Other arguments.

Details

This function calls [predict.HoltWinters](#) and constructs an object of class "forecast" from the results.

It is included for completeness, but the [ets](#) is recommended for use instead of [HoltWinters](#).

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.HoltWinters`.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals

level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

See Also

[predict.HoltWinters](#), [HoltWinters](#).

Examples

```
fit <- HoltWinters(WWWusage, gamma=FALSE)
plot(forecast(fit))
```

forecast.lm

Forecast a linear model with possible time series components

Description

forecast.lm is used to predict linear models, especially those involving trend and seasonality components.

Usage

```
## S3 method for class 'lm'
forecast(object, newdata, h=10, level=c(80,95), fan=FALSE,
         lambda=object$lambda, ...)
```

Arguments

object	Object of class "lm", usually the result of a call to lm or tslm .
newdata	An optional data frame in which to look for variables with which to predict. If omitted, it is assumed that the only variables are trend and season, and h forecasts are produced.
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to seq(50,99,by=1). This is suitable for fan plots.
h	Number of periods for forecasting. Ignored if newdata present.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.
...	Other arguments passed to predict.lm() .

Details

`forecast.lm` is largely a wrapper for `predict.lm()` except that it allows variables "trend" and "season" which are created on the fly from the time series characteristics of the data. Also, the output is reformatted into a forecast object.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.lm`.

An object of class "forecast" is a list containing at least the following elements:

<code>model</code>	A list containing information about the fitted model
<code>method</code>	The name of the forecasting method as a character string
<code>mean</code>	Point forecasts as a time series
<code>lower</code>	Lower limits for prediction intervals
<code>upper</code>	Upper limits for prediction intervals
<code>level</code>	The confidence values associated with the prediction intervals
<code>x</code>	The historical data for the response variable.
<code>residuals</code>	Residuals from the fitted model. That is x minus fitted values.
<code>fitted</code>	Fitted values

Author(s)

Rob J Hyndman

See Also

[tslm](#), [lm](#).

Examples

```
y <- ts(rnorm(120,0,3) + 1:120 + 20*sin(2*pi*(1:120)/12), frequency=12)
fit <- tslm(y ~ trend + season)
plot(forecast(fit, h=20))
```

forecast.stl

*Forecasting using stl objects***Description**

Returns forecasts obtained by either ETS or ARIMA models applied to the seasonally adjusted data from an STL decomposition.

Usage

```
## S3 method for class 'stl'
forecast(object, method=c("ets", "arima", "naive", "rwdrift"), etsmodel="ZZN",
         h=frequency(object$time.series)*2, level=c(80,95), fan=FALSE, lambda=NULL, ...)
stlf(x, h=frequency(x)*2, s.window=7, robust=FALSE, method=c("ets", "arima", "naive", "rwdrift"), etsmodel=c(80,95), fan=FALSE, lambda=NULL, ...)
```

Arguments

object	An object of class "stl". Usually the result of a call to stl .
x	A univariate numeric time series of class "ts"
method	Method to use for forecasting the seasonally adjusted series.
etsmodel	The ets model specification passed to ets . By default it allows any non-seasonal model. If <code>method!="ets"</code> , this argument is ignored.
h	Number of periods for forecasting.
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to <code>seq(50,99,by=1)</code> . This is suitable for fan plots.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, data transformed before model is estimated and back-transformed after forecasts are computed.
s.window	Either the character string "periodic" (default) or the span (in lags) of the loess window for seasonal extraction.
robust	If TRUE, robust fitting will used in the loess procedure within stl .
...	Other arguments passed to ets() or auto.arima() .

Details

`forecast.stl` seasonally adjusts the data from an STL decomposition, then uses either ETS or ARIMA models to forecast the result. The seasonal component from the last year of data is added back in to the forecasts. Note that the prediction intervals ignore the uncertainty associated with the seasonal component.

`stlf` takes a `ts` argument and applies a `stl` decomposition before calling `forecast.stl`.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.stl`.

An object of class "forecast" is a list containing at least the following elements:

<code>model</code>	A list containing information about the fitted model
<code>method</code>	The name of the forecasting method as a character string
<code>mean</code>	Point forecasts as a time series
<code>lower</code>	Lower limits for prediction intervals
<code>upper</code>	Upper limits for prediction intervals
<code>level</code>	The confidence values associated with the prediction intervals
<code>x</code>	The original time series (either object itself or the time series used to create the model stored as object).
<code>residuals</code>	Residuals from the fitted model. That is (possibly transformed) x minus fitted values.
<code>fitted</code>	Fitted values (one-step forecasts) on transformed scale if <code>lambda</code> is not NULL.

Author(s)

Rob J Hyndman

See Also

[forecast.ets](#), [forecast.Arima](#).

Examples

```
fit <- stl(USAccDeaths,s.window="periodic")
plot(forecast(fit))

plot(stlf(AirPassengers, lambda=BoxCox.lambda(AirPassengers)))
```

forecast.StructTS *Forecasting using Structural Time Series models*

Description

Returns forecasts and other information for univariate structural time series models.

Usage

```
## S3 method for class 'StructTS'
forecast(object, h=ifelse(object$coef["epsilon"] > 1e-10, 2*object$xtsp[3],10),
         level=c(80,95), fan=FALSE, lambda=NULL, ...)
```

Arguments

object	An object of class "StructTS". Usually the result of a call to StructTS .
h	Number of periods for forecasting
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to seq(50,99,by=1). This is suitable for fan plots.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.
...	Other arguments.

Details

This function calls `predict.StructTS` and constructs an object of class "forecast" from the results.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `forecast.StructTS`.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals

x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

See Also

[StructTS](#).

Examples

```
fit <- StructTS(WWWusage, "level")
plot(forecast(fit))
```

gas

Australian monthly gas production

Description

Australian monthly gas production: 1956–1995.

Usage

gas

Format

Time series data

Source

Australian Bureau of Statistics.

Examples

```
plot(gas)
seasonplot(gas)
tsdisplay(gas)
```

gold	<i>Daily morning gold prices</i>
------	----------------------------------

Description

Daily morning gold prices in US dollars. 1 January 1985 – 31 March 1989.

Usage

```
data(gold)
```

Format

Time series data

Source

Time Series Data Library. <http://robjhyndman.com/TSDL/>

Examples

```
tsdisplay(gold)
```

logLik.ets	<i>Log-Likelihood of an ets object</i>
------------	--

Description

Returns the log-likelihood of the ets model represented by object evaluated at the estimated parameters.

Usage

```
## S3 method for class 'ets'
logLik(object, ...)
```

Arguments

object	an object of class ets, representing an exponential smoothing state space model.
...	some methods for this generic require additional arguments. None are used in this method.

Value

the log-likelihood of the model represented by object evaluated at the estimated parameters.

Author(s)

Rob J Hyndman

References

Hyndman, R.J., Koehler, A.B., Ord, J.K., and Snyder, R.D. (2008) *Forecasting with exponential smoothing: the state space approach*, Springer-Verlag. <http://www.exponentialsmoothing.net>.

See Also

[ets](#)

Examples

```
fit <- ets(USAccDeaths)
logLik(fit)
```

ma	<i>Moving-average smoothing</i>
----	---------------------------------

Description

Computes a simple moving average smoother.

Usage

```
ma(x, order, centre=TRUE)
```

Arguments

x	Univariate time series
order	Order of moving average smoother
centre	If TRUE, then the moving average is centred.

Value

Numerical time series object containing the smoothed values.

Author(s)

Rob J Hyndman

See Also

[ksmooth](#), [decompose](#)

Examples

```
plot(wineind)
sm <- ma(wineind,order=12)
lines(sm,col="red")
```

 meanf

Mean Forecast

Description

Returns forecasts and prediction intervals for an iid model applied to x .

Usage

```
meanf(x, h=10, level=c(80,95), fan=FALSE, lambda=NULL)
```

Arguments

x	a numeric vector or time series
h	Number of periods for forecasting
level	Confidence levels for prediction intervals.
fan	If TRUE, level is set to seq(50,99,by=1). This is suitable for fan plots.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.

Details

The iid model is

$$Y_t = \mu + Z_t$$

where Z_t is a normal iid error. Forecasts are given by

$$Y_n(h) = \mu$$

where μ is estimated by the sample mean.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `meanf`.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
-------	--

method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

See Also

[rwf](#)

Examples

```
nile.fcast <- meanf(Nile, h=10)
plot(nile.fcast)
```

monthdays	<i>Number of days in each season</i>
-----------	--------------------------------------

Description

Returns number of days in each month or quarter of the observed time period.

Usage

```
monthdays(x)
```

Arguments

x time series

Details

Useful for month length adjustments

Value

Time series

Author(s)

Rob J Hyndman

Examples

```
par(mfrow=c(2,1))
plot(ldeaths,xlab="Year",ylab="pounds",
     main="Monthly deaths from lung disease (UK)")
ldeaths.adj <- ldeaths/monthdays(ldeaths)*365.25/12
plot(ldeaths.adj,xlab="Year",ylab="pounds",
     main="Adjusted monthly deaths from lung disease (UK)")
```

msts

Multi-Seasonal Time Series

Description

msts is an S3 class for multi seasonal time series objects, intended to be used for models that support multiple seasonal periods. The msts class inherits from the ts class and has an additional "msts" attribute which contains the vector of seasonal periods. All methods that work on a ts class, should also work on a msts class.

Usage

```
msts(data, seasonal.periods, ts.frequency=floor(max(seasonal.periods)), ... )
```

Arguments

data	A numeric vector, ts object, matrix or data frame. It is intended that the time series data is univariate, otherwise treated the same as ts().
seasonal.periods	A vector of the seasonal periods of the msts.
ts.frequency	The seasonal periods that should be used as frequency of the underlying ts object. The default value is max(seasonal.periods).
...	Arguments to be passed to the underlying call to ts(). For example start=c(1987,5).

Value

An object of class c("msts", "ts").

Author(s)

Slava Razbash and Rob J Hyndman

Examples

```
x <- msts(taylor, seasonal.periods=c(48,336), ts.frequency=48, start=2000+22/52)
y <- msts(USAccDeaths, seasonal.periods=12, ts.frequency=12, start=1949)
```

na.interp	<i>Interpolate missing values in a time series</i>
-----------	--

Description

Uses linear interpolation to replace missing values.

Usage

```
na.interp(x)
```

Arguments

x time series

Details

A more general and flexible approach is available using `na.approx` in the `zoo` package.

Value

Time series

Author(s)

Rob J Hyndman

Examples

```
data(gold)
plot(na.interp(gold))
```

naive	<i>Naive forecasts</i>
-------	------------------------

Description

`naive()` returns forecasts and prediction intervals for an ARIMA(0,1,0) random walk model applied to `x`. `snaive()` returns forecasts and prediction intervals from an ARIMA(0,0,0)(0,1,0)_m model where `m` is the seasonal period.

Usage

```
naive(x, h=10, level=c(80,95), fan=FALSE, lambda=NULL)
snaive(x, h=2*frequency(x), level=c(80,95), fan=FALSE, lambda=NULL)
```

Arguments

x	a numeric vector or time series
h	Number of periods for forecasting
level	Confidence levels for prediction intervals.
fan	If TRUE, level is set to seq(50,99,by=1). This is suitable for fan plots.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.

Details

These functions are simply convenient wrappers to [Arima](#) with the appropriate arguments to return naive and seasonal naive forecasts.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `naive` or `snaive`.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

See Also

[Arima](#), [rwf](#)

Examples

```
plot(naive(gold,h=50),include=200)
plot(snaive(wineind))
```

ndiffs	<i>Number of differences required for a stationary series</i>
--------	---

Description

Functions to estimate the number of differences required to make a given time series stationary. `ndiffs` estimates the number of first differences and `nsdiffs` estimates the number of seasonal differences.

Usage

```
ndiffs(x, alpha=0.05, test=c("kpss","adf", "pp"))
nsdiffs(x, m=frequency(x), test=c("ocsb","ch"))
```

Arguments

<code>x</code>	A univariate time series
<code>alpha</code>	Level of the test
<code>m</code>	Length of seasonal period
<code>test</code>	Type of unit root test to use

Details

`ndiffs` uses a unit root test to determine the number of differences required for time series `x` to be made stationary. If `test="kpss"`, the KPSS test is used with the null hypothesis that `x` has a stationary root against a unit-root alternative. Then the test returns the least number of differences required to pass the test at the level `alpha`. If `test="adf"`, the Augmented Dickey-Fuller test is used and if `test="pp"` the Phillips-Perron test is used. In both of these cases, the null hypothesis is that `x` has a unit root against a stationary root alternative. Then the test returns the least number of differences required to fail the test at the level `alpha`.

`nsdiffs` uses seasonal unit root tests to determine the number of seasonal differences required for time series `x` to be made stationary (possibly with some lag-one differencing as well). If `test="ch"`, the Canova-Hansen (1995) test is used (with null hypothesis of deterministic seasonality) and if `test="ocsb"`, the Osborn-Chui-Smith-Birchenhall (1988) test is used (with null hypothesis that a seasonal unit root exists).

Value

An integer.

Author(s)

Rob J Hyndman and Slava Razbash

References

- Canova F and Hansen BE (1995) "Are Seasonal Patterns Constant over Time? A Test for Seasonal Stability", *Journal of Business and Economic Statistics* **13**(3):237-252.
- Dickey DA and Fuller WA (1979), "Distribution of the Estimators for Autoregressive Time Series with a Unit Root", *Journal of the American Statistical Association* **74**:427-431.
- Kwiatkowski D, Phillips PCB, Schmidt P and Shin Y (1992) "Testing the Null Hypothesis of Stationarity against the Alternative of a Unit Root", *Journal of Econometrics* **54**:159-178.
- Osborn DR, Chui APL, Smith J, and Birchenhall CR (1988) "Seasonality and the order of integration for consumption", *Oxford Bulletin of Economics and Statistics* **50**(4):361-377.
- Osborn, D.R. (1990) "Seasonality and the order of integration in consumption", *International Journal of Forecasting*, **6**:327-336.
- Said E and Dickey DA (1984), "Testing for Unit Roots in Autoregressive Moving Average Models of Unknown Order", *Biometrika* **71**:599-607.

See Also

[auto.arima](#)

Examples

```
ndiffs(WWWusage)
nsdiffs(log(AirPassengers))
ndiffs(diff(log(AirPassengers),12))
```

plot.bats

Plot components from BATS model

Description

Produces a plot of the level, slope and seasonal components from a BATS or TBATS model.

Usage

```
## S3 method for class 'bats'
plot(x, main="Decomposition by BATS model", ...)
## S3 method for class 'tbats'
plot(x, main="Decomposition by TBATS model", ...)
```

Arguments

x	Object of class "ets".
main	Main title for plot.
...	Other plotting parameters passed to par .

Value

None. Function produces a plot

Author(s)

Rob J Hyndman

See Also

[bats](#), [tbats](#)

Examples

```
fit <- bats(USAccDeaths)
plot(fit)

fit2 <- tbats(USAccDeaths)
plot(fit2)
```

plot.ets

Plot components from ETS model

Description

Produces a plot of the level, slope and seasonal components from an ETS model.

Usage

```
## S3 method for class 'ets'
plot(x, ...)
```

Arguments

x Object of class “ets”.

... Other plotting parameters passed to [par](#).

Value

None. Function produces a plot

Author(s)

Rob J Hyndman

See Also

[ets](#)

Examples

```
fit <- ets(USAccDeaths)
plot(fit)
plot(fit, plot.type="single", ylab="", col=1:3)
```

plot.forecast

Forecast plot

Description

Plots historical data with forecasts and prediction intervals.

Usage

```
## S3 method for class 'forecast'
plot(x, include, plot.conf=TRUE, shaded=TRUE,
      shadebars=(length(x$mean)<5), shadecols=NULL, col=1, fcol=4,
      pi.col=1, pi.lty=2, ylim=NULL, main=NULL, ylab="", xlab="", type="l", ...)

## S3 method for class 'splineforecast'
plot(x, fitcol=2, type="o", pch=19, ...)
```

Arguments

x	Forecast object produced by forecast .
include	number of values from time series to include in plot
plot.conf	Logical flag indicating whether to plot prediction intervals.
shaded	Logical flag indicating whether prediction intervals should be shaded (TRUE) or lines (FALSE)
shadebars	Logical flag indicating if prediction intervals should be plotted as shaded bars (if TRUE) or a shaded polygon (if FALSE). Ignored if shaded=FALSE. Bars are plotted by default if there are fewer than five forecast horizons.
shadecols	Colors for shaded prediction intervals
col	the colour for the data line.
fcol	the colour for the forecast line.
pi.col	If shade=FALSE and plot.conf=TRUE, the prediction intervals are plotted in this colour.
pi.lty	If shade=FALSE and plot.conf=TRUE, the prediction intervals are plotted using this line type.
ylim	Limits on y-axis
main	Main title
ylab	Y-axis label
xlab	X-axis label

fitcol	Line colour for fitted values.
type	1-character string giving the type of plot desired. As for <code>plot.default</code> .
pch	Plotting character (if <code>type=="p"</code> or <code>type=="o"</code>).
...	additional arguments to <code>plot</code> .

Value

None.

Author(s)

Rob J Hyndman

References

Makridakis, Wheelwright and Hyndman (1998) *Forecasting: methods and applications*, Wiley: New York. <http://robjhyndman.com/forecasting/>.

See Also

[plot.ts](#)

Examples

```
deaths.fit <- hw(USAccDeaths,h=48)
plot(deaths.fit)
```

rwf	<i>Random Walk Forecast</i>
-----	-----------------------------

Description

Returns forecasts and prediction intervals for a random walk with drift model applied to `x`.

Usage

```
rwf(x, h=10, drift=FALSE, level=c(80,95), fan=FALSE, lambda=NULL)
```

Arguments

<code>x</code>	a numeric vector or time series
<code>h</code>	Number of periods for forecasting
<code>drift</code>	Logical flag. If TRUE, fits a random walk with drift model.
<code>level</code>	Confidence levels for prediction intervals.
<code>fan</code>	If TRUE, level is set to <code>seq(50,99,by=1)</code> . This is suitable for fan plots.
<code>lambda</code>	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.

Details

The random walk with drift model is

$$Y_t = c + Y_{t-1} + Z_t$$

where Z_t is a normal iid error. Forecasts are given by

$$Y_n(h) = ch + Y_n$$

. If there is no drift, the drift parameter $c=0$. Forecast standard errors allow for uncertainty in estimating the drift parameter.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `rwf`.

An object of class "forecast" is a list containing at least the following elements:

<code>model</code>	A list containing information about the fitted model
<code>method</code>	The name of the forecasting method as a character string
<code>mean</code>	Point forecasts as a time series
<code>lower</code>	Lower limits for prediction intervals
<code>upper</code>	Upper limits for prediction intervals
<code>level</code>	The confidence values associated with the prediction intervals
<code>x</code>	The original time series (either object itself or the time series used to create the model stored as object).
<code>residuals</code>	Residuals from the fitted model. That is <code>x</code> minus fitted values.
<code>fitted</code>	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

See Also

[arima](#), [meanf](#)

Examples

```
gold.fcast <- rwf(gold[1:60],h=50)
plot(gold.fcast)
```

seasadj	<i>Seasonal adjustment</i>
---------	----------------------------

Description

Returns seasonally adjusted data constructed by removing the seasonal component.

Usage

```
seasadj(object)
```

Arguments

object Object created by [decompose](#) or [stl](#).

Value

Univariate time series.

Author(s)

Rob J Hyndman

See Also

[stl](#), [decompose](#)

Examples

```
plot(AirPassengers)
lines(seasadj(decompose(AirPassengers,"multiplicative")),col=4)
```

seasonaldummy	<i>Seasonal dummy variables</i>
---------------	---------------------------------

Description

seasonaldummy and seasonaldummyf return matrices of dummy variables suitable for use in [arima](#), [lm](#) or [tslm](#). The last season is omitted and used as the control.

fourier and fourierf return matrices containing terms from a Fourier series, up to order K, suitable for use in [arima](#), [lm](#) or [tslm](#).

Usage

```
seasonaldummy(x)
seasonaldummyf(x,h)
fourier(x,K)
fourierf(x,K,h)
```

Arguments

x	Seasonal time series
h	Number of periods ahead to forecast
K	Maximum order of Fourier terms

Value

Numerical matrix with number of rows equal to the length(x) and number of columns equal to frequency(x)-1 (for seasonaldummy and seasonaldummyf or 2*K (for fourier or fourierf)).

Author(s)

Rob J Hyndman

Examples

```
plot(ldeaths)

# Using seasonal dummy variables
month <- seasonaldummy(ldeaths)
deaths.lm <- tslm(ldeaths ~ month)
tsdisplay(residuals(deaths.lm))
ldeaths.fcast <- forecast(deaths.lm,
  data.frame(month=I(seasonaldummyf(ldeaths,36))))
plot(ldeaths.fcast)

# A simpler approach to seasonal dummy variables
deaths.lm <- tslm(ldeaths ~ season)
ldeaths.fcast <- forecast(deaths.lm, h=36)
plot(ldeaths.fcast)

# Using Fourier series
X <- fourier(ldeaths,3)
deaths.lm <- tslm(ldeaths ~ X)
ldeaths.fcast <- forecast(deaths.lm,
  data.frame(X=I(fourierf(ldeaths,3,36))))
plot(ldeaths.fcast)
```

seasonplot	<i>Seasonal plot</i>
------------	----------------------

Description

Plots a seasonal plot as described in Makridakis, Wheelwright and Hyndman (1998, chapter 2).

Usage

```
seasonplot(x, s, season.labels=NULL, year.labels=FALSE,  
           year.labels.left=FALSE, type="o", main, ylab="",  
           xlab=NULL, col=1, labelgap=0.1, ...)
```

Arguments

x	a numeric vector or time series.
s	seasonal frequency of x
season.labels	Labels for each season in the "year"
year.labels	Logical flag indicating whether labels for each year of data should be plotted on the right.
year.labels.left	Logical flag indicating whether labels for each year of data should be plotted on the left.
type	plot type (as for <code>plot</code>)
main	Main title.
ylab	Y-axis label
xlab	X-axis label
col	Colour
labelgap	Distance between year labels and plotted lines
...	additional arguments to <code>plot</code> .

Value

None.

Author(s)

Rob J Hyndman

References

Makridakis, Wheelwright and Hyndman (1998) *Forecasting: methods and applications*, Wiley: New York. <http://robjhyndman.com/forecasting/>

See Also[monthplot](#)**Examples**

```
seasonplot(AirPassengers,col=rainbow(12),year.labels=TRUE)
```

ses

*Exponential smoothing forecasts***Description**

Returns forecasts and other information for exponential smoothing forecasts applied to x .

Usage

```
ses(x, h=10, level=c(80,95), fan=FALSE, ...)
holt(x, h=10, damped=FALSE, level=c(80,95), fan=FALSE, ...)
hw(x, h=2*frequency(x), seasonal="additive", damped=FALSE,
   level=c(80,95), fan=FALSE, ...)
```

Arguments

x	a numeric vector or time series
h	Number of periods for forecasting.
<code>damped</code>	If TRUE, use a damped trend.
<code>seasonal</code>	Type of seasonality in hw model. "additive" or "multiplicative"
<code>level</code>	Confidence level for prediction intervals.
<code>fan</code>	If TRUE, level is set to seq(50,99,by=1). This is suitable for fan plots.
<code>...</code>	Other arguments passed to ets.

Details

ses, holt and hw are simply convenient wrapper functions for forecast(ets(...)).

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by ets and associated functions.

An object of class "forecast" is a list containing at least the following elements:

<code>model</code>	A list containing information about the fitted model
--------------------	--

method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

References

Hyndman, R.J., Koehler, A.B., Snyder, R.D., Grose, S. (2002) "A state space framework for automatic forecasting using exponential smoothing methods", *International J. Forecasting*, **18**(3), 439–454.

Hyndman, R.J., Akram, Md., and Archibald, B. (2008) "The admissible parameter space for exponential smoothing models". *Annals of Statistical Mathematics*, **60**(2), 407–426.

See Also

[ets](#), [HoltWinters](#), [rwf](#), [arima](#).

Examples

```
fcast <- holt(airmiles)
plot(fcast)
deaths.fcast <- hw(USAccDeaths, h=48)
plot(deaths.fcast)
```

simulate.ets

Simulation from a time series model

Description

Returns a time series based on the model object object.

Usage

```
## S3 method for class 'ets'
simulate(object, nsim=length(object$x), seed=NULL, future=TRUE,
         bootstrap=FALSE, ...)
## S3 method for class 'ar'
simulate(object, nsim=object$n.used, seed=NULL, future=TRUE,
         bootstrap=FALSE, ...)
## S3 method for class 'Arima'
simulate(object, nsim=length(object$x), seed=NULL, xreg=NULL, future=TRUE,
         bootstrap=FALSE, ...)
## S3 method for class 'fracdiff'
simulate(object, nsim=object$n, seed=NULL, future=TRUE,
         bootstrap=FALSE, ...)
```

Arguments

object	An object of class "ets", "Arima" or "ar".
nsim	Number of periods for the simulated series
seed	Either NULL or an integer that will be used in a call to set.seed before simulating the time series. The default, NULL will not change the random generator state.
future	Produce sample paths that are future to and conditional on the data in object.
bootstrap	If TRUE, simulation uses resampled errors rather than normally distributed errors.
xreg	New values of xreg to be used for forecasting. Must have nsim rows.
...	Other arguments.

Value

An object of class "ts".

Author(s)

Rob J Hyndman

See Also

[ets](#), [Arima](#), [auto.arima](#), [ar](#), [arfima](#).

Examples

```
fit <- ets(USAccDeaths)
plot(USAccDeaths, xlim=c(1973,1982))
lines(simulate(fit, 36), col="red")
```

sindexf	<i>Forecast seasonal index</i>
---------	--------------------------------

Description

Returns vector containing the seasonal index for h future periods. If the seasonal index is non-periodic, it uses the last values of the index.

Usage

```
sindexf(object, h)
```

Arguments

object	Output from decompose or stl .
h	Number of periods ahead to forecast

Value

Time series

Author(s)

Rob J Hyndman

Examples

```
uk.stl <- stl(UKDriverDeaths,"periodic")
uk.sa <- seasadj(uk.stl)
uk.fcast <- holt(uk.sa,36)
seasf <- sindexf(uk.stl,36)
uk.fcast$mean <- uk.fcast$mean + seasf
uk.fcast$lower <- uk.fcast$lower + cbind(seasf,seasf)
uk.fcast$upper <- uk.fcast$upper + cbind(seasf,seasf)
uk.fcast$x <- UKDriverDeaths
plot(uk.fcast,main="Forecasts from Holt's method with seasonal adjustment")
```

splinef	<i>Cubic Spline Forecast</i>
---------	------------------------------

Description

Returns local linear forecasts and prediction intervals using cubic smoothing splines.

Usage

```
splinef(x, h=10, level=c(80,95), fan=FALSE, lambda=NULL)
```

Arguments

x	a numeric vector or time series
h	Number of periods for forecasting
level	Confidence level for prediction intervals.
fan	If TRUE, level is set to seq(50,99,by=1). This is suitable for fan plots.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, forecasts back-transformed via an inverse Box-Cox transformation.

Details

The cubic smoothing spline model is equivalent to an ARIMA(0,2,2) model but with a restricted parameter space. The advantage of the spline model over the full ARIMA model is that it provides a smooth historical trend as well as a linear forecast function. Hyndman, King, Pitrun, and Billah (2002) show that the forecast performance of the method is hardly affected by the restricted parameter space.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `meanf`.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

References

Hyndman, King, Pitrun and Billah (2005) Local linear forecasts using cubic smoothing splines. *Australian and New Zealand Journal of Statistics*, 47(1), 87-99. <http://robjhyndman.com/papers/splinefcast/>.

See Also

[smooth.spline](#), [arima](#), [holt](#).

Examples

```
fcast <- splines(uspop,h=5)
plot(fcast)
summary(fcast)
```

subset.ts	<i>Subsetting a time series</i>
-----------	---------------------------------

Description

The main purpose of this function is to extract the values of a specific season in each year. For example, to extract all values for the month of May from a time series.

Usage

```
## S3 method for class 'ts'
subset(x, subset=NULL, month=NULL, quarter=NULL, season=NULL, ...)
```

Arguments

x	a univariate time series to be subsetted
subset	optional logical expression indicating elements to keep; missing values are taken as false.
month	Character list of months to retain. Partial matching on month names used.
quarter	Numeric list of quarters to retain.
season	Numeric list of seasons to retain.
...	Other arguments, unused.

Value

If one season per year is extracted, then a ts object is returned with frequency 1. Otherwise, a numeric vector is returned with no ts attributes.

Author(s)

Rob J Hyndman

See Also

[subset](#)

Examples

```
plot(subset(gas,month="November"))
subset(woolyrnq,quarter=3)
```

taylor	<i>Half-hourly electricity demand</i>
--------	---------------------------------------

Description

Half-hourly electricity demand in England and Wales from Monday 5 June 2000 to Sunday 27 August 2000. Discussed in Taylor (2003), and kindly provided by James W Taylor.

Usage

```
taylor
```

Format

Time series data

Source

James W Taylor

References

Taylor, J.W. (2003) Short-term electricity demand forecasting using double seasonal exponential smoothing. *Journal of the Operational Research Society*, **54**, 799-805.

Examples

```
plot(taylor)
```

tbats	<i>TBATS model (Exponential smoothing state space model with Box-Cox transformation, ARMA errors, Trend and Seasonal components)</i>
-------	--

Description

Fits a TBATS model applied to y , as described in De Livera, Hyndman & Snyder (2011). Parallel processing is used by default to speed up the computations.

Usage

```
tbats(y, use.box.cox=NULL, use.trend=NULL, use.damped.trend=NULL,
      seasonal.periods=NULL, use.arma.errors=TRUE, use.parallel=TRUE, num.cores=NULL, ...)
```

Arguments

<code>y</code>	The time series to be forecast. Can be numeric, msts or ts. Only univariate time series are supported.
<code>use.box.cox</code>	TRUE/FALSE indicates whether to use the Box-Cox transformation or not. If NULL then both are tried and the best fit is selected by AIC.
<code>use.trend</code>	TRUE/FALSE indicates whether to include a trend or not. If NULL then both are tried and the best fit is selected by AIC.
<code>use.damped.trend</code>	TRUE/FALSE indicates whether to include a damping parameter in the trend or not. If NULL then both are tried and the best fit is selected by AIC.
<code>seasonal.periods</code>	If <code>y</code> is numeric then seasonal periods can be specified with this parameter.
<code>use.arma.errors</code>	TRUE/FALSE indicates whether to include ARMA errors or not. If TRUE the best fit is selected by AIC. If FALSE then the selection algorithm does not consider ARMA errors.
<code>use.parallel</code>	TRUE/FALSE indicates whether or not to use parallel processing.
<code>num.cores</code>	The number of parallel processes to be used if using parallel processing. If NULL then the number of logical cores is detected.
<code>...</code>	Additional parameters to be passed to <code>auto.arima</code> when choose an ARMA(p, q) model for the errors.

Value

An object with class `c("tbats", "bats")`. The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `bats` and associated functions.

Author(s)

Slava Razbash and Rob J Hyndman

References

De Livera, A.M., Hyndman, R.J., & Snyder, R. D. (2011), Forecasting time series with complex seasonal patterns using exponential smoothing, *Journal of the American Statistical Association*, **106**(496), 1513-1527.

Examples

```
fit <- tbats(USAccDeaths)
plot(forecast(fit))
## Not run:
taylor.fit <- tbats(taylor)
plot(forecast(taylor.fit))

## End(Not run)
```

thetaf	<i>Theta method forecast</i>
--------	------------------------------

Description

Returns forecasts and prediction intervals for a theta method forecast.

Usage

```
thetaf(x, h=10, level=c(80,95), fan=FALSE)
```

Arguments

x	a numeric vector or time series
h	Number of periods for forecasting
level	Confidence levels for prediction intervals.
fan	If TRUE, level is set to seq(50,99,by=1). This is suitable for fan plots.

Details

The theta method of Assimakopoulos and Nikolopoulos (2000) is equivalent to simple exponential smoothing with drift. This is demonstrated in Hyndman and Billah (2003). Prediction intervals are computed using the underlying state space model.

Value

An object of class "forecast".

The function `summary` is used to obtain and print a summary of the results, while the function `plot` produces a plot of the forecasts and prediction intervals.

The generic accessor functions `fitted.values` and `residuals` extract useful features of the value returned by `rwf`.

An object of class "forecast" is a list containing at least the following elements:

model	A list containing information about the fitted model
method	The name of the forecasting method as a character string
mean	Point forecasts as a time series
lower	Lower limits for prediction intervals
upper	Upper limits for prediction intervals
level	The confidence values associated with the prediction intervals
x	The original time series (either object itself or the time series used to create the model stored as object).
residuals	Residuals from the fitted model. That is x minus fitted values.
fitted	Fitted values (one-step forecasts)

Author(s)

Rob J Hyndman

References

Assimakopoulos, V. and Nikolopoulos, K. (2000). The theta model: a decomposition approach to forecasting. *International Journal of Forecasting* **16**, 521-530.

Hyndman, R.J., and Billah, B. (2003) Unmasking the Theta method. *International J. Forecasting*, **19**, 287-290.

See Also

[arima](#), [meanf](#), [rwf](#), [ses](#)

Examples

```
nile.fcast <- thetaf(Nile)
plot(nile.fcast)
```

 tsdisplay

Time series display

Description

Plots a time series along with its acf and either its pacf, lagged scatterplot or spectrum.

Usage

```
tsdisplay(x, plot.type="partial", points=TRUE, ci.type="white",
  lag.max, na.action=na.interp,
  main=NULL, ylab="", xlab="", pch=1, cex=0.5, ...)
```

Arguments

x	a numeric vector or time series.
plot.type	type of plot to include in lower right corner. Possible values are "partial", "scatter" or "spectrum".
points	logical flag indicating whether to show the individual points or not in the time plot.
ci.type	type of confidence limits for ACF. Possible values are as for acf .
lag.max	the maximum lag to plot for the acf and pacf. A suitable value is selected by default if the argument is missing.
na.action	how to handle missing values. Default is to use linear interpolation.
main	Main title.
ylab	Y-axis label

xlab	X-axis label
pch	Plotting character
cex	Character size
...	additional arguments to acf .

Value

None.

Author(s)

Rob J Hyndman

References

Makridakis, Wheelwright and Hyndman (1998) *Forecasting: methods and applications*, Wiley: New York. <http://robjhyndman.com/forecasting/>

See Also

[plot.ts](#), [acf](#)

Examples

```
tsdisplay(diff(WWWusage))
```

tslm

Fit a linear model with time series components

Description

tslm is used to fit linear models to time series including trend and seasonality components.

Usage

```
tslm(formula, data, lambda=NULL, ...)
```

Arguments

formula	an object of class "formula" (or one that can be coerced to that class): a symbolic description of the model to be fitted.
data	an optional data frame, list or environment (or object coercible by <code>as.data.frame</code> to a data frame) containing the variables in the model. If not found in data, the variables are taken from <code>environment(formula)</code> , typically the environment from which <code>lm</code> is called.
lambda	Box-Cox transformation parameter. Ignored if NULL. Otherwise, data are transformed via a Box-Cox transformation.
...	Other arguments passed to <code>lm()</code> .

Details

`tslm` is largely a wrapper for `lm()` except that it allows variables "trend" and "season" which are created on the fly from the time series characteristics of the data. The variable "trend" is a simple time trend and "season" is a factor indicating the season (e.g., the month or the quarter depending on the frequency of the data).

Value

Returns an object of class "lm".

Author(s)

Rob J Hyndman

See Also

[forecast.lm](#), [lm](#).

Examples

```
y <- ts(rnorm(120,0,3) + 1:120 + 20*sin(2*pi*(1:120)/12), frequency=12)
fit <- tslm(y ~ trend + season)
plot(forecast(fit, h=20))
```

wineind

Australian total wine sales

Description

Australian total wine sales by wine makers in bottles \leq 1 litre. Jan 1980 – Aug 1994.

Usage

```
wineind
```

Format

Time series data

Source

Time Series Data Library. <http://robjhyndman.com/TSDL/>

Examples

```
tsdisplay(wineind)
```

woolyrnq

Quarterly production of woollen yarn in Australia

Description

Quarterly production of woollen yarn in Australia: tonnes. Mar 1965 – Sep 1994.

Usage

woolyrnq

Format

Time series data

Source

Time Series Data Library. <http://robjhyndman.com/TSDL/>

Examples

```
tsdisplay(woolyrnq)
```

Index

*Topic **datasets**

gas, 35
gold, 36
taylor, 58
wineind, 63
woolryrnq, 64

*Topic **hplot**

plot.bats, 44
plot.ets, 45

*Topic **htest**

dm.test, 17

*Topic **models**

CV, 16

*Topic **stats**

forecast.lm, 30
tslm, 62

*Topic **ts**

accuracy, 3
Acf, 4
arfima, 5
Arima, 6
arima.errors, 8
auto.arima, 9
bats, 11
BoxCox, 12
BoxCox.lambda, 13
croston, 14
dm.test, 17
dshw, 18
ets, 20
fitted.Arima, 22
forecast, 23
forecast.Arima, 24
forecast.bats, 26
forecast.ets, 27
forecast.HoltWinters, 29
forecast.stl, 32
forecast.StructTS, 34
logLik.ets, 36

ma, 37
meanf, 38
monthdays, 39
msts, 40
na.interp, 41
naive, 41
ndiffs, 43
plot.forecast, 46
rwf, 47
seasadj, 49
seasonaldummy, 49
seasonplot, 51
ses, 52
simulate.ets, 53
sindexf, 55
splinef, 55
subset.ts, 57
tbats, 58
thetaf, 60
tsdisplay, 61

accuracy, 3
Acf, 4
acf, 4, 5, 61, 62
AIC, 16
ar, 24, 25, 54
arfima, 5, 24, 25, 54
Arima, 6, 11, 25, 42, 54
arima, 5–10, 22–25, 48, 49, 53, 57, 61
arima.errors, 8
auto.arima, 5, 6, 9, 24, 25, 32, 44, 54

bats, 11, 26, 27, 45
best.arima (auto.arima), 9
BoxCox, 12, 14
BoxCox.lambda, 13, 13

croston, 14, 24
CV, 16

decompose, 37, 49, 55

- dm.test, 17
- dshw, 18
- ets, 19, 20, 23, 27–29, 32, 37, 45, 53, 54
- fitted.Arima, 22
- forecast, 23, 46
- forecast.ar (forecast.Arima), 24
- forecast.Arima, 22, 23, 24, 24, 33
- forecast.bats, 26
- forecast.ets, 23, 24, 27, 27, 33
- forecast.fracdiff, 6, 25
- forecast.fracdiff (forecast.Arima), 24
- forecast.HoltWinters, 24, 29
- forecast.lm, 30, 63
- forecast.stl, 32
- forecast.StructTS, 24, 34
- forecast.tbats (forecast.bats), 26
- forecast.ts, 23
- fourier (seasonaldummy), 49
- fourierf (seasonaldummy), 49
- fracdiff, 5, 6, 24
- gas, 35
- gold, 36
- holt, 24, 57
- holt (ses), 52
- HoltWinters, 19, 22, 29, 30, 53
- hw, 24
- hw (ses), 52
- InvBoxCox (BoxCox), 12
- ksmooth, 37
- lm, 16, 30, 31, 49, 62, 63
- logLik.ets, 36
- ma, 37
- meanf, 24, 38, 48, 61
- monthdays, 39
- monthplot, 52
- msts, 40
- na.interp, 41
- naive, 41
- ndiffs, 10, 43
- nsdiffs, 10
- nsdiffs (ndiffs), 43
- Pacf (Acf), 4
- par, 44, 45
- plot, 47, 51
- plot.bats, 44
- plot.default, 47
- plot.ets, 45
- plot.forecast, 46
- plot.splineforecast (plot.forecast), 46
- plot.tbats (plot.bats), 44
- plot.ts, 47, 62
- predict.ar, 25
- predict.Arima, 25
- predict.HoltWinters, 29, 30
- predict.lm, 30, 31
- print.forecast (forecast), 23
- residuals, 9
- rwf, 22, 24, 39, 42, 47, 53, 61
- seasadj, 49
- seasonaldummy, 49
- seasonaldummyf (seasonaldummy), 49
- seasonplot, 51
- ses, 15, 24, 52, 61
- set.seed, 54
- simulate.ar (simulate.ets), 53
- simulate.Arima (simulate.ets), 53
- simulate.ets, 53
- simulate.fracdiff (simulate.ets), 53
- sindexf, 55
- smooth.spline, 57
- snaive (naive), 41
- splinef, 24, 55
- stl, 32, 49, 55
- stl (forecast.stl), 32
- stlf, 23
- stlf (forecast.stl), 32
- StructTS, 34, 35
- subset, 57
- subset.ts, 57
- summary.forecast (forecast), 23
- taylor, 58
- tbats, 27, 45, 58
- thetaf, 24, 60
- tsdisplay, 61
- tslm, 16, 30, 31, 49, 62
- wineind, 63
- woolyrnq, 64